

Monte Carlo Search Algorithm Discovery for One Player Games

Francis Maes, David Lupien St-Pierre and Damien Ernst

Abstract—Much current research in AI and games is being devoted to Monte Carlo search (MCS) algorithms. While the quest for a single unified MCS algorithm that would perform well on all problems is of major interest for AI, practitioners often know in advance the problem they want to solve, and spend plenty of time exploiting this knowledge to customize their MCS algorithm in a problem-driven way. We propose an MCS algorithm discovery scheme to perform this in an automatic and reproducible way. We first introduce a grammar over MCS algorithms that enables inducing a rich space of candidate algorithms. Afterwards, we search in this space for the algorithm that performs best on average for a given distribution of training problems. We rely on multi-armed bandits to approximately solve this optimization problem. The experiments, generated on three different domains, show that our approach enables discovering algorithms that outperform several well known MCS algorithms such as Upper Confidence bounds applied to Trees and Nested Monte Carlo search. We also show that the discovered algorithms are generally quite robust with respect to changes in the distribution over the training problems.

Index Terms—Monte-Carlo Search, Algorithm Selection, Grammar of Algorithms

I. INTRODUCTION

MONTE CARLO search (MCS) algorithms rely on random simulations to evaluate the quality of states or actions in sequential decision making problems. Most of the recent progress in MCS algorithms has been obtained by integrating smart procedures to select the simulations to be performed. This has led to, among other things, the Upper Confidence bounds applied to Trees algorithm (UCT, [1]) that was popularized thanks to breakthrough results in computer Go [2]. This algorithm relies on a game tree to store simulation statistics and uses this tree to bias the selection of future simulations. While UCT is one way to combine random simulations with tree search techniques, many other approaches are possible. For example, the Nested Monte Carlo (NMC) search algorithm [3], which obtained excellent results in the last General Game Playing competition¹ [4], relies on nested levels of search and does not require storing a game tree.

How to best bias the choice of simulations is still an active topic in MCS-related research. Both UCT and NMC are attempts to provide generic techniques that perform well on a wide range of problems and that work with little or no prior knowledge. While working on such generic algorithms is definitely relevant to AI, MCS algorithms are in practice

widely used in a totally different scenario, in which a significant amount of prior knowledge is available about the game or the sequential decision making problem to be solved.

People applying MCS techniques typically spend plenty of time exploiting their knowledge of the target problem so as to design more efficient problem-tailored variants of MCS. Among the many ways to do this, one very common practice is automatic hyper-parameter tuning. By way of example, the parameter $C > 0$ of UCT is in nearly all applications tuned through a more or less automated trial and error procedure. While hyper-parameter tuning is a very simple form of problem-driven algorithm selection, most of the advanced algorithm selection work is done by humans, i.e., by researchers that modify or invent new algorithms to take the specificities of their problem into account.

The comparison and development of new MCS algorithms given a target problem is mostly a manual search process that takes much human time and is error prone. Thanks to modern computing power, automatic discovery is becoming a credible approach for partly automating this process. In order to investigate this research direction, we focus on the simplest case of (fully observable) deterministic single-player games. Our contribution is twofold. First, we introduce a grammar over the MCS algorithms that enables generating a very rich space of MCS algorithms. It also describes several well known MCS algorithms, using a particularly compact and elegant description. Second, we assume that there is access to a distribution over the training problems that reflects this prior knowledge, and adopt as a performance criterion for an algorithm its mean performance on problems drawn from this distribution. We rely on multi-armed bandits for identifying in a computational efficient way the algorithm that maximizes this criterion. Our approach is tested on three different domains. The results show that it enables discovering new variants of MCS that significantly outperform generic algorithms such as UCT or NMC on each of these domains. We further show the good robustness properties of the discovered algorithms by slightly changing the characteristics of the problem.

This paper is structured as follows. Section II formalizes the class of sequential decision making problems considered in this paper and formalizes the corresponding MCS algorithm discovery problem. Section III describes our grammar over MCS algorithms and describes several well known MCS algorithms in terms of this grammar. Section IV formalizes the search for a good MCS algorithm as a multi-armed bandit problem. We experimentally evaluate our approach on different domains in Section V. Finally, we discuss related work in

Francis Maes, David Lupien St-Pierre and Damien Ernst are at the Department of Electrical Engineering and Computer Science, University of Liège, Liège, Belgium.

E-mail: {francis.maes, dlspierre, dernst}@ulg.ac.be

¹<http://games.stanford.edu>

Section VI and conclude in Section VII.

II. PROBLEM STATEMENT

We consider the class of finite horizon fully observable deterministic sequential decision making problems. A problem P is a triple (x_1, f, g) where $x_1 \in \mathcal{X}$ is the initial state, f is the transition function, and g is the reward function. The dynamics of a problem is described by

$$x_{t+1} = f(x_t, u_t) \quad t = 1, 2, \dots, T, \quad (1)$$

where for all t , the state x_t is an element of the state space \mathcal{X} and the action u_t is an element of the action space \mathcal{U} . In the context of one player games, x_t denotes the current state of the game and \mathcal{U} are the possible moves. We make no assumptions on the nature of \mathcal{X} but assume that \mathcal{U} is finite. We denote by \mathcal{U}_x the set of actions (moves) which are legal in state $x \in \mathcal{X}$. We assume that when starting from x_1 , the system enters a final state after T steps and we denote by $\mathcal{F} \subset \mathcal{X}$ the set of these final states². A final state $x \in \mathcal{F}$ is associated to rewards $g(x) \in \mathbb{R}$ that should be maximized.

A search algorithm $A(\cdot)$ is a stochastic algorithm that explores the possible sequences of actions to approximatively maximize

$$A(P = (x_1, f, g)) \simeq \underset{u_1, \dots, u_T}{\operatorname{argmax}} g(x_{T+1}), \quad (2)$$

subject to $x_{t+1} = f(x_t, u_t)$ and $u_t \in \mathcal{U}_{x_t}$. In order to fulfill this task, the algorithm is given a finite amount of computational time, referred to as the *budget*. Moreover, to facilitate its reproducibility, we focus in this paper on a budget expressed as the maximum number $B > 0$ of sequences (u_1, \dots, u_T) that can be evaluated, or, equivalently, as the number of calls to the reward function $g(\cdot)$. Note, however, that it is trivial in our approach to replace this definition by other budget measures, such as CPU time.

We express our knowledge of the problem as a distribution over problem instances \mathcal{D}_P , from which we can sample any number of random instances $P \sim \mathcal{D}_P$. The quality of a search algorithm $A^B(\cdot)$ with budget B on these instances is denoted by $J_A^B(\mathcal{D}_P)$ and is defined as the expected quality of solutions found on problems drawn from \mathcal{D}_P :

$$J_A^B(\mathcal{D}_P) = \mathbb{E}_{P \sim \mathcal{D}_P} \{ \mathbb{E}_{x_{T+1} \sim A^B(P)} \{ g(x_{T+1}) \} \}, \quad (3)$$

where $x_{T+1} \sim A^B(P)$ denotes the final states returned by algorithm A with budget B on problem P .

Given a class of candidate algorithms \mathcal{A} and given the budget B , the algorithm discovery problem amounts to selecting an algorithm $A^* \in \mathcal{A}$ of maximal quality:

$$A^* = \underset{A \in \mathcal{A}}{\operatorname{argmax}} J_A^B(\mathcal{D}_P). \quad (4)$$

²In many problems, the time at which the game enters a final state is not fixed, but depends on the actions played so far. It should however be noted that it is possible to make these problems fit this fixed finite time formalism by postponing artificially the end of the game until T . This can be done, for example, by considering that when the game ends before T , a “pseudo final state” is reached from which, whatever the actions taken, the game will reach the real final state in T .

The two main contributions of this paper are: (i) a grammar that enables inducing a rich space \mathcal{A} of candidate MCS algorithms, and (ii) an efficient procedure to approximately solve Eq. 4.

III. A GRAMMAR FOR MONTE-CARLO SEARCH ALGORITHMS

All MCS algorithms share some common underlying general principles: random simulations, look-ahead search, time-receding control, and bandit-based selection. The grammar that we introduce in this section aims at capturing these principles in a pure and atomic way. We first give an overall view of our approach, then present in detail the components of our grammar, and finally describe previously proposed algorithms by using this grammar.

A. Overall view

We call *search components* the elements on which our grammar operates. Formally, a search component is a stochastic algorithm that, when given a partial sequence of actions (u_1, \dots, u_{t-1}) , generates one or multiple completions (u_t, \dots, u_T) and evaluates them using the reward function $g(\cdot)$. The search components are denoted by $S \in \mathcal{S}$, where \mathcal{S} is the space of all possible search components. Let S be a particular search component. We define the search algorithm $A_S \in \mathcal{A}$ as the algorithm that, given the problem P , executes S repeatedly with an empty partial sequence of actions $()$, until the computational budget is exhausted. The search algorithm A_S then returns the sequence of actions (u_1, \dots, u_T) that led to the highest reward $g(\cdot)$.

In order to generate a rich class of search components—hence a rich class of search algorithms—in an inductive way, we rely on search-component *generators*. Such generators are functions $\Psi : \Theta \rightarrow \mathcal{S}$ that define a search component $S = \Psi(\theta) \in \mathcal{S}$ when given a set of parameters $\theta \in \Theta$. Our grammar is composed of five search component generators: $\Psi \in \{\text{simulate}, \text{repeat}, \text{lookahead}, \text{step}, \text{select}\}$. Here is a short description of the search components they generate:

- *simulate*(π^{simu}): Generate one single completion (u_t, \dots, u_T) using simulation policy π^{simu} .
- *repeat*(N, S): Execute N times the search component S .
- *lookahead*(S): For each valid action u_t , move to successor state $f(x_t, u_t)$ and run sub-search component S .
- *step*(S): For each $\tau \in [t, T]$, run sub-search component S , select action u_τ towards the best currently found final state, and move one step ahead.
- *select*(π^{sel}, S): Use a selection policy recursively to select one or multiple actions u_t, u_{t+1}, \dots , run search component S , and back propagate the sub-search result.

Note that four of our search component generators are parametrized by sub-search components. For example, *step* and *lookahead* are functions $\mathcal{S} \rightarrow \mathcal{S}$. These functions can be nested recursively to generate more and more evolved search components. We construct the space of search algorithms \mathcal{A} by performing this in a systematic way, as detailed in Section IV-A.

B. Search components

Figure 1 describes our five search component generators. Note that we distinguish between search component *inputs* and search component generator *parameters*. All our search components have the same two inputs: the sequence of already decided actions (u_1, \dots, u_{t-1}) and the current state $x_t \in \mathcal{X}$. The parameters differ from one search component generator to another. For example, *simulate* is parametrized by a simulation policy π^{simu} and *repeat* is parametrized by the number of repetitions $N > 0$ and by a sub-search component. We now give a detailed description of these search component generators.

Simulate The *simulate* generator is parametrized by a policy $\pi^{simu} \in \Pi^{simu}$ which is a stochastic mapping from states to actions: $u \sim \pi^{simu}(x)$. In order to generate the completion (u_t, \dots, u_T) , *simulate* repeatedly samples actions u_τ according to $\pi^{simu}(x_\tau)$ and performs transitions $x_{\tau+1} = f(x_\tau, u_\tau)$ until reaching a final state. A default choice for the simulation policy is the uniformly random policy, defined as

$$\mathbb{E}\{\pi^{random}(x) = u\} = \begin{cases} \frac{1}{|\mathcal{U}_x|} & \text{if } u \in \mathcal{U}_x \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

Once the completion (u_t, \dots, u_T) is fulfilled, the whole sequence (u_1, \dots, u_T) is *yielded*. This operation is detailed in Figure 2 and proceeds as follows: (i) it computes the reward of the final state x_{T+1} , (ii) if the reward is larger than the largest reward found previously, it replaces the best current solution, and (iii) if the budget B is exhausted, it stops the search.

Since algorithm A_P repeats P until the budget is exhausted, the search algorithm $A_{simulate} \in \mathcal{A}$ is *Iterative Sampling*: it samples B random trajectories (u_1, \dots, u_T) , evaluates each of the final state rewards $g(x_{T+1})$, and returns the best found final state. Note that, in the YIELD procedure, the variables relative to the best current solution (r^* and (u_1^*, \dots, u_T^*)) are defined locally for each search component, whereas the *numCalls* counter is global to the search algorithm. This means that if S is a search component composed of different nested levels of search (see the examples below), the best current solution is kept in memory at each level of search.

Repeat Given a positive integer $N > 0$ and a search component $S \in \mathcal{S}$, *repeat*(N, S) is the search component that repeats N times the search component S . For example, $S = \text{repeat}(10, \text{simulate}(\pi^{simu}))$ is the search component that draws 10 random simulations using π^{simu} . The corresponding search algorithm A_S is again iterative sampling, since search algorithms repeat their search component until the budget is exhausted. In Figure 1, we use the INVOKE operation each time a search component calls a sub-search component. This operation is detailed in Figure 2 and ensures that no sub-search algorithm is called when a final state is reached, i.e., when $t = T + 1$.

Look-ahead For each legal move $u_t \in \mathcal{U}_{x_t}$, *lookahead*(S) computes the successor state $x_{t+1} = f(x_t, u_t)$ and runs the sub-search component $S \in \mathcal{S}$ starting from the sequence (u_1, \dots, u_t) . For example, *lookahead*(*simulate*(π^{simu})) is the search component that, given the partial sequence

Fig. 1. Search component generators

SIMULATE((u_1, \dots, u_{t-1}), x_t)
Param: $\pi^{simu} \in \Pi^{simu}$

for $\tau = t$ to T **do**
 $u_\tau \sim \pi^{simu}(x_\tau)$
 $x_{\tau+1} \leftarrow f(x_\tau, u_\tau)$
end for
YIELD((u_1, \dots, u_T))

REPEAT((u_1, \dots, u_{t-1}), x_t)
Param: $N > 0, S \in \mathcal{S}$

for $i = 1$ to N **do**
INVOKE($S, (u_1, \dots, u_{t-1}), x_t$)
end for

LOOKAHEAD((u_1, \dots, u_{t-1}), x_t)
Param: $S \in \mathcal{S}$

for $u_t \in \mathcal{U}_{x_t}$ **do**
 $x_{t+1} \leftarrow f(x_t, u_t)$
INVOKE($S, (u_1, \dots, u_t), x_{t+1}$)
end for

STEP((u_1, \dots, u_{t-1}), x_t)
Param: $S \in \mathcal{S}$

for $\tau = t$ to T **do**
INVOKE($S, (u_1, \dots, u_{\tau-1}), x_\tau$)
 $u_\tau \leftarrow u_\tau^*$
 $x_{\tau+1} \leftarrow f(x_\tau, u_\tau)$
end for

SELECT((u_1, \dots, u_{t-1}), x_t)
Param: $\pi^{sel} \in \Pi^{sel}, S \in \mathcal{S}$

for $\tau = t$ to T **do** ▷ Select
 $u_\tau \sim \pi^{sel}(x)$
 $x_{\tau+1} \leftarrow f(x_\tau, u_\tau)$
if $n(x_{\tau+1}) = 0$ **then**
break
end if
end for
 $t_{leaf} \leftarrow \tau$

INVOKE($S, (u_1, \dots, u_{t_{leaf}}), x_{t_{leaf}+1}$) ▷ Sub-search

for $\tau = t_{leaf}$ to 1 **do** ▷ Backpropagate
 $n(x_{\tau+1}) \leftarrow n(x_{\tau+1}) + 1$
 $n(x_\tau, u_\tau) \leftarrow n(x_\tau, u_\tau) + 1$
 $s(x_\tau, u_\tau) \leftarrow s(x_\tau, u_\tau) + r^*$
end for
 $n(x_1) \leftarrow n(x_1) + 1$

(u_1, \dots, u_{t-1}) , generates one random trajectory for each legal next action $u_t \in \mathcal{U}_{x_t}$. Multiple-step look-

Fig. 2. Yield and invoke commands

Require: $g : \mathcal{F} \rightarrow \mathbb{R}$, the reward function
Require: $B > 0$, the computational budget
Initialize global: $numCalls \leftarrow 0$
Initialize local: $r^* \leftarrow -\infty$
Initialize local: $(u_1^*, \dots, u_T^*) \leftarrow \emptyset$
procedure YIELD((u_1, \dots, u_T))
 $r = g(x)$
 if $r > r^*$ **then**
 $r^* \leftarrow r$
 $(u_1^*, \dots, u_T^*) \leftarrow (u_1, \dots, u_T)$
 end if
 $numCalls \leftarrow numCalls + 1$
 if $numCalls = B$ **then**
 stop search
 end if
end procedure
procedure INVOKE($S \in \mathcal{S}, (u_1, \dots, u_{t-1}) \in \mathcal{U}^*, x_t \in \mathcal{X}$)
 if $t \leq T$ **then**
 $S((u_1, \dots, u_{t-1}), x_t)$
 else
 yield (u_1, \dots, u_T)
 end if
end procedure

ahead search strategies naturally write themselves with nested calls to *lookahead*. As an example, $lookahead(lookahead(lookahead(simulate(\pi^{simu}))))$ is a search component that runs one random trajectory per legal combination of the three next actions (u_t, u_{t+1}, u_{t+2}) .

Step For each remaining time step $\tau \in [t, T]$, $step(S)$ runs the sub-search component S , extracts the action u_τ from (u_1^*, \dots, u_T^*) (the best currently found action sequence, see Figure 2), and performs transition $x_{\tau+1} = f(x_\tau, u_\tau)$. The search component generator *step* enables implementing time receding search mechanisms, e.g., $step(repeat(100, simulate(\pi^{simu})))$ is the search component that selects the actions (u_1, \dots, u_T) one by one, using 100 random trajectories to select each action. As a more evolved example, $step(lookahead(lookahead(repeat(10, simulation(\pi^{simu}))))$ is a time receding strategy that performs 10 random simulations for each two first actions (u_t, u_{t+1}) to decide which action u_t to select.

Select This search component generator implements most of the behaviour of a Monte Carlo Tree Search (MCTS, [1]). It relies on a game tree, which is a non-uniform look-ahead tree with nodes corresponding to states and edges corresponding to transitions. The role of this tree is twofold: it stores statistics on the outcomes of sub-searches and it is used to bias sub-searches towards promising sequences of actions. A search component $select(\pi^{sel}, S)$ proceeds in three steps: the *selection* step relies on the statistics stored in the game tree to select a (typically small) sub-sequence of actions $(u_t, \dots, u_{t_{leaf}})$, the *sub-search* step invokes the sub-search component $S \in \mathcal{S}$ starting from $(u_1, \dots, u_{t_{leaf}})$, and the

backpropagation step updates the statistics to take into account the sub-search result.

We use the following notation to denote the information stored by the look-ahead tree: $n(x, u)$ is the number of times the action u was selected in state x , $s(x, u)$ is the sum of rewards that were obtained when running *sub-search* after having selected action u in state x , and $n(x)$ is the number of times state x was selected: $n(x) = \sum_{u \in \mathcal{U}_x} n(x, u)$. In order to quantify the quality of a sub-search, we rely on the reward of the best solution that was tried during that sub-search: $r^* = \max g(x)$. In the simplest case, when the sub-search component is $S = simulate(\pi^{simu})$, r^* is the reward associated to the final state obtained by making the random simulation with policy π^{simu} , as usual in MCTS. In order to select the first actions, *selection* relies on a *selection policy* $\pi^{sel} \in \Pi^{sel}$, which is a stochastic function that, when given all stored information related to state x (i.e., $n(x)$, $n(x, u)$, and $s(x, u)$, $\forall u \in \mathcal{U}_x$), selects an action $u \in \mathcal{U}_x$. The selection policy has two contradictory goals to pursue: *exploration*, trying new sequences of actions to increase knowledge, and *exploitation*, using current knowledge to bias computational efforts towards promising sequences of actions. Such exploration/exploitation dilemmas are usually formalized as a multi-armed bandit problem, hence π^{sel} is typically one of policies commonly found in the multi-armed bandit literature. The probably most well known such policy is UCB-1 [5]:

$$\pi_C^{ucb-1}(x) = \operatorname{argmax}_{u \in \mathcal{U}_x} \frac{s(x, u)}{n(x, u)} + C \sqrt{\frac{\ln n(x, u)}{n(x)}}, \quad (6)$$

where division by zero returns $+\infty$ and where $C > 0$ is a hyper-parameter that enables the control of the exploration / exploitation tradeoff.

C. Description of previously proposed algorithms

Our grammar enables generating a very large class of MCS algorithms, which includes several already proposed algorithms. We now overview these algorithms, which can be described particularly compactly and elegantly thanks to our grammar:

- The simplest Monte Carlo algorithm is *Iterative Sampling*. This algorithm draws random simulations until the computational time is elapsed and returns the best solution found [6]:

$$is = simulate(\pi^{simu}). \quad (7)$$

- In general, iterative sampling is used during a certain time to decide which action to select (or which move to play) at each step of the decision problem. The corresponding search component is

$$is' = step(repeat(N, simulate(\pi^{simu}))), \quad (8)$$

where N is the number of simulations performed for each decision step.

- The *Reflexive Monte Carlo* search algorithm introduced in [7] proposes using a Monte Carlo search of a given level

to improve the search of the upper level. The proposed algorithm can be described as follows:

$$rmc(N_1, N_2) = \text{step}(\text{repeat}(N_1, \text{step}(\text{repeat}(N_2, \text{simulate}(\pi^{\text{simu}}))))), \quad (9)$$

where N_1 and N_2 are called the number of meta-games and the number of games, respectively.

- The Nested Monte Carlo (NMC) search algorithm is a recursively defined algorithm generalizing the ideas of Reflexive Monte Carlo search. NMC can be described in a very natural way by our grammar. The basic search level $l = 0$ of NMC simply performs a random simulation:

$$nmc(0) = \text{simulate}(\pi^{\text{random}}). \quad (10)$$

The level $l > 0$ of NMC relies on level $l - 1$ in the following way:

$$nmc(l) = \text{step}(\text{lookahead}(nmc(l - 1))). \quad (11)$$

- Single-player MCTS selects actions one after the other. In order to select one action, it relies on *select* combined with random simulations. The corresponding search component is thus

$$mcts(\pi^{\text{sel}}, \pi^{\text{simu}}, N) = \text{step}(\text{repeat}(N, \text{select}(\pi^{\text{sel}}, \text{simulate}(\pi^{\text{simu}})))), \quad (12)$$

where N is the number of iterations allocated to each decision step. UCT is one of the best known variants of MCTS. It relies on the $\pi_C^{\text{ucb}-1}$ selection policy and is generally used with a uniformly random simulation policy:

$$uct(C, N) = mcts(\pi_C^{\text{ucb}-1}, \pi^{\text{random}}, N). \quad (13)$$

- In the spirit of the work on nested Monte Carlo, the authors of [8] proposed the Meta MCTS approach, which replaces the simulation part of an upper-level MCTS algorithm by a whole lower-level MCTS algorithm. While they presented this approach in the context of two-player games, we can describe its equivalent for one-player games with our grammar:

$$\text{metamcts}(\pi^{\text{sel}}, \pi^{\text{simu}}, N_1, N_2) = \text{step}(\text{repeat}(N_1, \text{select}(\pi^{\text{sel}}, mcts(\pi^{\text{sel}}, \pi^{\text{simu}}, N_2)))). \quad (14)$$

where N_1 and N_2 are the budgets for the higher-level and lower-level MCTS algorithms, respectively.

In addition to offering a framework for describing these already proposed algorithms, our grammar enables generating a huge number of new hybrid MCS variants. We give, in the next section, a procedure to automatically identify the best such variant for a given problem.

IV. BANDIT-BASED ALGORITHM DISCOVERY

We now move to the problem of solving Eq. 4, i.e., of finding, for a given problem, the best algorithm A from among a large class \mathcal{A} of algorithms derived with the grammar previously defined. Solving this algorithm discovery problem exactly is impossible in the general case since the objective function involves two infinite expectations: one over the problems $P \sim \mathcal{D}_P$ and another over the outcomes of the algorithm. In order to approximately solve Eq. 4, we adopt the formalism of multi-armed bandits and proceed in two steps: we first construct a finite set of candidate algorithms $\mathcal{A}_{D,\Gamma} \subset \mathcal{A}$ (Section IV-A), and then treat each of these algorithms as an arm and use a multi-armed bandit policy to select how to allocate computational time to the performance estimation of the different algorithms (Section IV-B). It is worth mentioning that this two-step approach follows a general methodology for automatic discovery that we already successfully applied to multi-armed bandit policy discovery [9], [10], reinforcement learning policy discovery [11], and optimal control policy discovery [12].

A. Construction of the algorithm space

We measure the complexity of a search component $S \in \mathcal{S}$ using its *depth*, defined as the number of nested search components constituting S , and denote this quantity by $\text{depth}(S)$. For example, $\text{depth}(\text{simulate}(\pi^{\text{simu}}))$ is 1, $\text{depth}(\text{uct})$ is 4, and $\text{depth}(nmc(3))$ is 7.

Note that *simulate*, *repeat*, and *select* have parameters which are not search components: the simulation policy π^{simu} , the number of repetitions N , and the selection policy π^{sel} , respectively. In order to generate a finite set of algorithms using our grammar, we rely on predefined finite sets of possible values for each of these parameters. We denote by Γ the set of these finite domains. The discrete set $\mathcal{A}_{D,\Gamma}$ is constructed by enumerating all possible algorithms up to depth D with constants Γ , and is pruned using the following rules:

- *Canonization of repeat*: Both search components $S_1 = \text{step}(\text{repeat}(2, \text{repeat}(5, S_{\text{sub}})))$ and $S_2 = \text{step}(\text{repeat}(5, \text{repeat}(2, S_{\text{sub}})))$ involve running S_{sub} 10 times at each step. In order to avoid having this kind of algorithm duplicated, we collapse nested *repeat* components into single *repeat* components. With this rule, S_1 and S_2 both reduce to $\text{step}(\text{repeat}(10, S_{\text{sub}}))$.
- *Removal of nested selects*: A search component such as $\text{select}(\pi^{\text{sel}}, \text{select}(\pi^{\text{sel}}, S))$ is ill-defined, since the inner *select* will be called with a different initial state x_t each time, making it behave randomly. We therefore exclude search components involving two directly nested *selects*.
- *Removal of repeat-as-root*: Remember that the MCS algorithm $A_S \in \mathcal{A}$ runs S repeatedly until the computational budget is exhausted. Due to this repetition, algorithms such as $A_{\text{simulate}(\pi^{\text{simu}})}$ and $A_{\text{repeat}(10, \text{simulate}(\pi^{\text{simu}}))}$ are equivalent. To remove these duplicates, we reject all search components whose “root” is *repeat*.

In the following, ν denote the cardinality of the set of candidate algorithms: $\mathcal{A}_{D,\Gamma} = \{A_1, \dots, A_\nu\}$. To illustrate the construction of this set, consider a simple case where the

Depth 1-2	Depth 3	
simlate	lookahead(repeat(2, sim))	step(repeat(2, sim))
	lookahead(repeat(10, sim))	step(repeat(10, sim))
lookahead(sim)	lookahead(lookahead(sim))	step(lookahead(sim))
step(sim)	lookahead(step(sim))	step(step(sim))
select(sim)	lookahead(select(sim))	step(select(sim))
	select(repeat(2, sim))	select(repeat(10, sim))
	select(lookahead(sim))	select(step(sim))

TABLE I
UNIQUE ALGORITHMS UP TO DEPTH 3

maximum depth is $D = 3$ and where the constants Γ are $\pi^{simu} = \pi^{random}$, $N \in \{2, 10\}$, and $\pi^{sel} = \pi_C^{ucb-1}$. The corresponding space $\mathcal{A}_{D,\Gamma}$ contains $\nu = 18$ algorithms. These algorithms are given in Table I, where we use *sim* as an abbreviation for *simulate*(π^{simu}).

B. Bandit-based algorithm discovery

One simple approach to approximately solve Eq. 4 is to estimate the objective function through an empirical mean computed using a finite set of training problems $\{P^{(1)}, \dots, P^{(M)}\}$, drawn from \mathcal{D}_P :

$$J_A^B(\mathcal{D}_P) \simeq \frac{1}{M} \sum_{i=1}^M g(x_{T+1}) | x_{T+1} \sim A^B(P^{(i)}), \quad (15)$$

where x_{T+1} denotes one outcome of algorithm A with budget B on problem $P^{(i)}$. To solve Eq. 4, one can then compute this approximated objective function for all algorithms $A \in \mathcal{A}_{D,\Gamma}$ and simply return the algorithm with the highest score. While extremely simple to implement, such an approach often requires an excessively large number of samples M to work well, since the variance of $g(\cdot)$ may be quite large.

In order to optimize Eq. 4 in a smarter way, we propose to formalize this problem as a multi-armed bandit problem. To each algorithm $A_k \in \mathcal{A}_{D,\Gamma}$, we associate an arm. Pulling the arm k for the t_k th time involves selecting the problem $P^{(t_k)}$ and running the algorithm A_k once on this problem. This leads to a reward associated to arm k whose value is the reward $g(x_{T+1})$ that comes with the solution x_{T+1} found by algorithm A_k . The purpose of multi-armed bandit algorithms is to process the sequence of observed rewards to select in a smart way the next algorithm to be tried, so that when the time allocated to algorithm discovery is exhausted, one (or several) high-quality algorithm(s) can be identified. How to select arms so as to identify the best one in a finite amount of time is known as the *pure exploration* multi-armed bandit problem [13]. It has been shown that index based policies based on upper confidence bounds such as UCB-1 were also good policies for solving pure exploration bandit problems. Our optimization procedure works thus by repeatedly playing arms according to such a policy. In our experiments, we perform a fixed number of such iterations. In practice, this multi-armed bandit approach can provide an answer at anytime, but returning the algorithm A_k with the currently highest empirical reward mean.

C. Discussion

Note that other approaches could be considered for solving our algorithm discovery problem. In particular, optimization over expression spaces induced by a grammar such as ours is often solved using Genetic Programming (GP, [14]). GP works by evolving a population of solutions, which, in our case, would be MCS algorithms. At each iteration, the current population is evaluated, the less good solutions are removed, and the best solutions are used to construct new candidates using mutation and cross-over operations. Most existing GP algorithms assume that the objective function is (at least approximately) deterministic. One major advantage of the bandit-based approach is to natively take into account the stochasticity of the objective function and its decomposability into problems. Thanks to the bandit formulation, badly performing algorithms are quickly rejected and the computational power is more and more focused on the most promising algorithms.

The main strengths of our bandit-based approach are the following. First, it is simple to implement and does not require entering into the details of complex mutation and cross-over operators. Second, it has only one hyper-parameter (the exploration/exploitation coefficient), for which there exists a robust default setting. Finally, since it is based on exhaustive search and on multi-armed bandit theory, formal guarantees can easily be derived to bound the regret, i.e., the difference between the performance of the best algorithm and the performance of the algorithm discovered [13], [15], [16].

Our approach is restricted to relatively small depths D since it relies on exhaustive search. In our case, we believe that many interesting MCS algorithms can be described using search components with low depth. In our experiments, we used $D = 5$, which already provides many original hybrid algorithms that deserve further research. If this limit was too restrictive, a major way of improvement would consist in combining the idea of bandits with the ideas of GP. In this spirit, the authors of [17] recently proposed a hybrid approach in which the selection of the members of a new population is posed as a multi-armed bandit problem. This enables combining the best of the two approaches: multi-armed bandits enable taking natively into account the stochasticity and decomposability of the objective function, while GP cross-over and mutation operators are used to generate new candidates dynamically in a smart way.

V. EXPERIMENTS

We now apply our automatic algorithm discovery approach to three different testbeds: Sudoku, Symbolic Regression, and Morpion Solitaire. The aim of our experiments was to show that our approach discovers MCS algorithms that significantly outperform several generic (problem independent) MCS algorithms: outperforms them on the training instances, on new testing instances, and even on instances drawn from distributions different from the original distribution used for the learning.

We first describe the experimental protocol in Section V-A. Sections V-B, V-C, and V-D then give the results obtained

on the three domains. Finally, Section V-E gives an overall discussion of our results.

A. Protocol

We first describe the two different kinds of algorithms that we compare in our experiments: generic algorithms and discovered algorithms.

Generic algorithms The generic algorithms are Nested Monte Carlo, Upper Confidence bounds applied to Trees, Look-ahead Search, and Iterative sampling. The search components for Nested Monte Carlo (*nmc*), UCT (*uct*), and Iterative sampling (*is*) have already been defined in Section III-C. The search component for Look-ahead Search of level $l > 0$ is defined by $la(l) = \text{step}(larec(l))$, where

$$larec(l) = \begin{cases} \text{lookahead}(larec(l-1)) & \text{if } l > 0 \\ \text{simulate}(\pi^{\text{random}}) & \text{otherwise.} \end{cases} \quad (16)$$

For both $la(\cdot)$ and $nmc(\cdot)$, we try all values within the range $[1, 5]$ for the level parameter. Note that $la(1)$ and $nmc(1)$ are equivalent, since both are defined by the search component $\text{step}(\text{lookahead}(\text{simulate}(\pi^{\text{random}})))$. For $uct(\cdot)$, we try the following values of C : $\{0, 0.3, 0.5, 1.0\}$ and set the budget per step to $\frac{B}{T}$, where B is the total budget and T is the horizon of the problem. This leads to the following set of generic algorithms: $\{nmc(2), nmc(3), nmc(4), nmc(5), is, la(1), la(2), la(3), la(4), la(5), uct(0), uct(0.3), uct(0.5), \text{and } uct(1)\}$. Note that we omit the $\frac{B}{T}$ parameter in uct for the sake of conciseness.

Discovered algorithms In order to generate the set of candidate algorithms, we used the following constants Γ : *repeat* can be used with 2, 5, 10, or 100 repetitions; and *select* relies on the UCB1 selection policy from Eq. (6) with the constants $\{0, 0.3, 0.5, 1.0\}$. We create a pool of algorithms by exhaustively generating all possible combinations of the search components up to depth $D = 5$. We apply the pruning rules described in Section IV-A, which results in a set of $\nu = 3, 155$ candidate MCS algorithms.

Algorithm discovery In order to carry out the algorithm discovery, we used a UCB policy for $100 \times \nu$ time steps, i.e., each candidate algorithm was executed 100 times on average. As discussed in Section IV-B, each bandit step involves running one of the candidate algorithms on a problem $P \sim \mathcal{D}_P$. We refer to \mathcal{D}_P as the *training distribution* in the following. Once we have played the UCB policy for $100 \times \nu$ time steps, we sort the algorithms by their average training performance and report the ten best algorithms.

Evaluation Since algorithm discovery is a form of “learning from examples”, care must be taken with overfitting issues. Indeed, the discovered algorithms may perform well on the training problems P while performing poorly on other problems drawn from \mathcal{D}_P . Therefore, to evaluate the MCS algorithms, we used a set of 1,000 *testing problems* $P \sim \mathcal{D}_P$ which are different from the training problems. We then evaluate the score of an algorithm as the mean performance obtained when running it once on each testing problem.

In each domain, we further test the algorithms either by changing the budget B and/or by using a new distribution \mathcal{D}'_P that differs from the training distribution \mathcal{D}_P . In each such experiment, we draw 100 problems from \mathcal{D}'_P and run the algorithm once on each problem.

In one domain (Morpion Solitaire), we used a particular case of our general setting, in which there was a single training problem P , i.e., the distribution \mathcal{D}_P was degenerate and always returned the same P . In this case, we focused our analysis on the robustness of the discovered algorithms when tested on a new problem P' and/or with a new budget B .

B. Sudoku

Sudoku, a Japanese term meaning “singular number”, is a popular puzzle played around the world. The Sudoku puzzle is made of a grid of $G^2 \times G^2$ cells, which is structured into blocks of size $G \times G$. When starting the puzzle, some cells are already filled in and the objective is to fill in the remaining cells with the numbers 1 through G^2 so that

- no row contains two instances of the same number,
- no column contains two instances of the same number,
- no block contains two instances of the same number.

Sudoku is of particular interest in our case because each Sudoku grid corresponds to a different initial state x_1 . Thus, a good algorithm $A(\cdot)$ is one that intrinsically has the versatility to face a wide variety of Sudoku grids.

In our implementation, we maintain for each cell the list of numbers that could be put in that cell without violating any of the three previous rules. If one of these lists becomes empty then the grid cannot be solved and we pass to a final state (see Footnote 2). Otherwise, we select the subset of cells whose number-list has the lowest cardinality, and define one action $u \in \mathcal{U}_x$ per possible number in each of these cells (as in [3]). The reward associated to a final state is its proportion of filled cells, hence a reward of 1 is associated to a perfectly filled grid.

We sample the initial states x_1 by filling 33% randomly selected cells as proposed in [3]. We denote by $\text{Sudoku}(G)$ the distribution over Sudoku problems obtained with this procedure (in the case of $G^2 \times G^2$ games). Even though Sudoku is most usually played with $G = 3$ [18], we carry out the algorithm discovery with $G = 4$ to make the problem more difficult. Our training distribution was thus $\mathcal{D}_P = \text{Sudoku}(4)$ and we used a training budget of $B = 1,000$. To evaluate the performance and robustness of the algorithms found, we tested the MCS algorithms on two distributions: $\mathcal{D}_P = \text{Sudoku}(4)$ and $\mathcal{D}'_P = \text{Sudoku}(5)$, using a budget of $B = 1,000$.

Table II presents the results, where the scores are the average number of filled cells, which is given by the reward times the total number of cells G^4 . Note that the results have been sorted by decreasing testing scores on \mathcal{D}_P . In each column, we underline both the best generic algorithm and the best discovered algorithm and show in bold all cases in which a discovered algorithm outperforms all tested generic algorithms.

The first result we observe is the overall better performance of the discovered algorithms over the generic ones. When

TABLE II
RANKING AND ROBUSTNESS OF ALGORITHMS DISCOVERED WHEN APPLIED TO SUDOKU

Name	Search Component	Rank	Sudoku(4)	Sudoku(5)
Dis#1	select(step(repeat(select(simulate, 1), 5)), 1)	1	197.3	480.7
Dis#8	step(select(repeat(select(simulate, 0.5), 5), 0))	2	197.2	481.2
Dis#2	step(repeat(step(repeat(simulate, 5)), 10))	3	197.2	484.9
Dis#7	lookahead(step(repeat(select(simulate, 0.3), 5)))	4	197.0	487.5
Dis#6	step(step(repeat(select(simulate, 0), 5)))	5	196.8	482.0
Dis#4	step(step(step(select(simulate, 0.5))))	6	196.7	491.3
Dis#10	select(step(repeat(select(simulate, 0.3), 5))	7	196.5	480.3
Dis#9	lookahead(step(step(select(simulate, 1))))	8	196.3	493.6
Dis#5	select(step(repeat(simulate, 5)), 0.5)	9	196.0	478.8
Dis#3	step(select(step(simulate), 1))	10	195.9	496.5
uct(0.5)		11	194.7	481.6
uct(0)		12	194.5	<u>485.5</u>
la(1)		13	193.7	429.7
nmc(2)		14	192.9	423.4
uct(0.3)		15	192.8	480.3
nmc(3)		16	192.4	428.6
uct(1)		17	191.0	482.3
nmc(4)		18	190.3	424.50
nmc(5)		19	190.0	425.60
la(2)		20	173.5	391.7
la(3)		21	168.7	389.4
la(4)		22	168.5	388.9
la(5)		23	168.4	386.4
is		24	168.1	389.40

testing on the distribution $\mathcal{D}_P = \text{Sudoku}(4)$, we observe that all ten discovered algorithms outperform the best generic algorithm, $uct(0.5)$. This shows that our grammar generated new algorithms which are specifically relevant to the $\text{Sudoku}(4)$ task. When running the algorithms on the $\text{Sudoku}(5)$ games, we observe that $uct(0)$ becomes the best generic algorithm. In this case, four of our discovered algorithms still remain better than UCT and all ten discovered algorithms still perform well. $\text{Sudoku}(4)$ and $\text{Sudoku}(5)$ are thus sufficiently close problems to make algorithms discovered for $\text{Sudoku}(4)$ relevant to $\text{Sudoku}(5)$.

We observe some frequent patterns in the search components of the discovered algorithms: five search components are based on a double nested *step* and nine out of the ten discovered algorithms rely on *select*.

C. Real Valued Symbolic Regression

Symbolic Regression consists in searching in a large space of symbolic expressions for the one that best fits a given regression dataset. Usually this problem is treated using Genetic Programming approaches. In the line of [19], we here consider MCS techniques as an interesting alternative to Genetic Programming. In order to apply MCS techniques, we encode the expressions as sequences of symbols. We adopt the Reverse Polish Notation (RPN) to avoid the use of parentheses. As an example, the sequence $[a, b, +, c, *]$ encodes the expression $(a+b)*c$. The alphabet of symbols we used is $\{x, 1, +, -, *, /, \sin, \cos, \log, \exp, stop\}$. The initial state x_1 is the empty RPN sequence. Each action u then adds one of these symbols to the sequence. When computing the set of valid actions \mathcal{U}_x , we reject symbols that lead to invalid RPN sequences, such as $[+, +, +]$. A final state is reached either when the sequence length is equal to a predefined maximum

Target Expression $f^P(\cdot)$	Domain
$x^3 + x^2 + x$	$[-1, 1]$
$x^4 + x^3 + x^2 + x$	$[-1, 1]$
$x^5 + x^4 + x^3 + x^2 + x$	$[-1, 1]$
$x^6 + x^5 + x^4 + x^3 + x^2 + x$	$[-1, 1]$
$\sin(x^2) \cos(x) - 1$	$[-1, 1]$
$\sin(x) + \sin(x + x^2)$	$[-1, 1]$
$\log(x + 1) + \log(x^2 + 1)$	$[0, 2]$
\sqrt{x}	$[0, 4]$

TABLE III
SYMBOLIC REGRESSION TESTBED: TARGET EXPRESSIONS AND DOMAINS.

Target Expression $f^P(\cdot)$	Domain
$x^3 - x^2 - x$	$[-1, 1]$
$x^4 - x^3 - x^2 - x$	$[-1, 1]$
$x^4 + \sin(x)$	$[-1, 1]$
$\cos(x^3) + \sin(x + 1)$	$[-1, 1]$
$\sqrt{x} + x^2$	$[0, 4]$
$x^6 + 1$	$[-1, 1]$
$\sin(x^3 + x^2)$	$[-1, 1]$
$\log(x^3 + 1) + x$	$[0, 2]$

TABLE IV
SYMBOLIC REGRESSION ROBUSTNESS TESTBED: TARGET EXPRESSIONS AND DOMAINS.

T or when the symbol *stop* is played. In our experiments, we performed the training with a maximal length of $T = 11$. The reward associated to a final state is equal to $1 - mae$, where mae is the mean absolute error associated to the expression built.

We used a synthetic benchmark, which is classical in the field of Genetic Programming [20]. To each problem P of this benchmark is associated a target expression $f^P(\cdot) \in \mathbb{R}$, and

TABLE V
RANKING AND ROBUSTNESS OF THE ALGORITHMS DISCOVERED WHEN APPLIED TO SYMBOLIC REGRESSION

Name	Search Component	Rank	$T = 11$	$T = 21$	$T = 11, B = 10^5$	\mathcal{D}'_P
Dis#1	step(step(lookahead(lookahead(simulate))))	1	0.066	0.080	0.036	0.101
Dis#5	step(repeat(lookahead(lookahead(simulate)), 2))	2	0.069	0.085	0.037	0.106
Dis#2	step(lookahead(lookahead(repeat(simulate, 2))))	3	0.070	0.086	0.038	0.100
Dis#7	step(lookahead(lookahead(select(simulate, 1))))	4	0.070	0.087	0.040	0.103
Dis#6	step(lookahead(lookahead(select(simulate, 0))))	5	0.071	0.088	0.040	0.110
Dis#4	step(lookahead(select(lookahead(simulate), 0)))	5	0.071	0.087	0.039	0.101
Dis#3	step(lookahead(lookahead(simulate)))	5	0.071	0.086	0.056	0.100
la(2)		5	<u>0.071</u>	<u>0.086</u>	0.056	<u>0.100</u>
Dis#10	step(lookahead(select(lookahead(simulate), 0.3)))	9	0.072	0.093	0.040	0.108
Dis#8	step(lookahead(repeat(lookahead(simulate), 2)))	9	0.072	0.085	0.040	0.112
la(3)		11	0.073	0.094	<u>0.053</u>	0.101
Dis#9	step(repeat(select(lookahead(simulate), 0.3), 5))	12	0.078	0.093	0.048	0.099
nmc(2)		13	0.081	0.103	0.054	0.109
nmc(3)		14	0.083	0.105	<u>0.053</u>	0.118
la(4)		15	0.089	0.115	0.057	0.101
nmc(4)		16	0.094	0.106	0.059	0.140
la(1)		17	0.098	0.114	0.066	0.120
la(5)		18	0.099	0.124	0.058	0.101
is		19	0.120	0.142	0.086	0.138
nmc(5)		20	0.121	0.122	0.069	0.140
uct(0)		21	0.151	0.130	0.124	0.185
uct(1)		22	0.154	0.127	0.118	0.160
uct(0.3)		23	0.155	0.129	0.135	0.177
uct(0.5)		24	0.156	0.127	0.124	0.184

the aim is to re-discover this target expression given a finite set of samples $(x, f^P(x))$. Table III illustrates these target expressions. In each case, we used 20 samples $(x, f^P(x))$, where x was obtained by taking uniformly spaced elements from the indicated domains. The training distribution \mathcal{D}_P was the uniform distribution over the eight problems given in Table III.

The training budget was $B = 10,000$. We evaluate the robustness of the algorithms found in three different ways: by changing the maximal length T from 11 to 21, by increasing the budget B from 10,000 to 100,000 and by testing them on another distribution of problems \mathcal{D}'_P . The distribution \mathcal{D}'_P is the uniform distribution over the eight new problems given in Table IV.

The results are shown in Table V, where we report directly the *mae* scores (lower is better). As in the Sudoku domain, we observe the generally strong performance of the discovered algorithms. There are two generic algorithms, $la(2)$ and $la(3)$, ranked 5th and 11th, respectively, that perform better than some of the discovered algorithms. We note, however, that $la(2)$ has been rediscovered (Dis#3) and that most discovered algorithms exhibit a look-ahead search structure: $step(lookahead(lookahead(\cdot)))$.

When setting the maximal length to $T = 21$, the discovered algorithms still outperform all the generic algorithms except $la(2)$. When increasing the testing budget to $B = 100,000$, $la(3)$ becomes more interesting than $la(2)$ and nine discovered algorithms out of the ten outperform $la(3)$. These results thus show that the algorithms discovered by our approach are robust both w.r.t. to the maximal length T and to the budget B .

In our last experiment with the distribution \mathcal{D}'_P , we again observe that all ten discovered algorithms behave particularly well by outperforming all the generic algorithms except $la(2)$ and $la(3)$. This result is particularly interesting since it shows

that our approach was enable to discover algorithms that work well for symbolic regression in general, not only for some particular problems.

When looking at the discovered search components, we observe that *lookahead* is intensively used in all good MCS algorithms. This is a strong indication that look-ahead search is a good strategy for this kind of expression discovery problems.

D. Morpion Solitaire

The classic game of morpion solitaire [21] is a single player, pencil and paper game, whose world record has been improved several times over the past few years using MCS techniques [3], [7], [22]. This game is illustrated in Figure 3. The initial state x_1 is an empty cross of points drawn on the intersections of the grid. Each action places a new point at a grid intersection in such a way that it forms a new line segment connecting consecutive points that include the new one. New lines can be drawn horizontally, vertically, and diagonally. The game is over when no further actions can be taken. The goal of the game is to maximize the number of lines drawn before the game ends, hence the reward associated to final states is this number³.

There exist two variants of the game: “Disjoint” and “Touching”. “Touching” allows parallel lines to share an endpoint, whereas “Disjoint” does not. Line segments with different directions are always permitted to share points. The game is NP-hard [23] and presumed to be infinite under certain configurations. In this paper, we treat the $5D$ and $5T$ versions of the game, where 5 is the number of consecutive points to form a line, D means disjoint, and T means touching.

³In practice, we normalize this reward by dividing it by 100 to make it approximately fit into the range $[0, 1]$. Thanks to this normalization, we can keep using the same constants for both the UCB policy used in the algorithm discovery and the UCB policy used in *select*.




Fig. 3. A random policy that plays the game Morpion Solitaire 5T: initial grid; after 1 move; after 10 moves; game end.

We performed the algorithm discovery in a “single training problem” scenario: the training distribution \mathcal{D}_P always returns the same problem P , corresponding to the 5T version of the game. The initial state of P was the one given in the leftmost part of Figure 3. The training budget was set to $B = 10,000$. To evaluate the robustness of the algorithms, we, on the one hand, evaluated them on the 5D variant of the problem and, on the other hand, changed the evaluation budget from 10,000 to 100,000. The former provides a partial answer to how rule-dependent these algorithms are, while the latter gives insight into the impact of the budget on the algorithms’ ranking.

The results of our experiments on Morpion Solitaire are given in Table VI. As before, the grammar was rich enough to generate several MCS algorithms which outperformed all the tested generic algorithms, with average scores of up to 91.43. Among the generic algorithms, Nested Monte Carlo gave the best results (90.68), which is 0.33 below the worst of the ten discovered algorithms.

When moving to the 5D rules, we observe that all ten discovered algorithms still outperform the best generic algorithm. This is particularly impressive, since it is known that the structure of good solutions strongly differs between the 5D and 5T versions of the game [21]. The last column of Table VI gives the performance of the algorithms with budget $B = 10^5$. We observe that all ten discovered algorithms also outperform the best generic algorithm in this case. Furthermore, the increase in the budget seems to also increase the gap between the discovered and the generic algorithms.

When looking at the discovered search components, we observe the same kinds of patterns as for Sudoku: double nested *step*, and the use of *select*.

E. Discussion

We have seen that on each of our three testbeds, all ten discovered algorithms outperform the best generic algorithm. This clearly demonstrates that our approach is able to generate new MCS algorithms specifically tailored to the given class of problems. We have performed a study of the robustness of these algorithms by either changing the problem distribution or by varying the budget B , and found that the algorithms discovered can outperform generic algorithms even on problems

significantly different from those used for the training.

The importance of each component of the grammar depends heavily on the problem. For instance, in Symbolic Regression, all ten best algorithms discovered rely on two nested *lookahead* components, whereas in Sudoku and Morpion, *step* and *select* appear in the majority of the best algorithms discovered.

VI. RELATED WORK

Methods for automatically discovering MCS algorithms can be characterized through three main components: the space of candidate algorithms, the performance criterion, and the search method for finding the best element in the space of candidate algorithms.

Usually, researchers consider spaces of candidate algorithms that only differ in the values of their constants. In such a context, the problem amounts to tuning the constants of a generic MCS algorithm. Most of the research related to the tuning of these constants takes as performance criterion the mean score of the algorithm over the distribution of the problem instances. Many search algorithms have been proposed for computing the best constants. For instance, [24] employs a grid search approach combined with self-playing, [25] uses cross-entropy as a search method to tune an agent playing GO, [26] presents a generic black-box optimization method based on local quadratic regression, [27] uses Estimation Distribution Algorithms with Gaussian distributions, [28] uses Thompson Sampling, and [29] uses, as in the present paper, a multi-armed bandit approach. The paper [30] studies the influence of the tuning of MCS algorithms on their asymptotic consistency and shows that pathological behaviour may occur with tuning. It also proposes a tuning method to avoid such behaviour.

Research papers that have reported empirical evaluations of several MCS algorithms in order to find the best one are also related to this automatic discovery problem. The space of candidate algorithms in such cases is the set of algorithms they compare, and the search method is an exhaustive search procedure. As a few examples, [24] reports on a comparison between algorithms that differ in their selection policy, [31] and [32] compare improvements of the UCT algorithm (RAVE and progressive bias) with the original one on the game of GO,

TABLE VI
RANKING AND ROBUSTNESS OF ALGORITHMS DISCOVERED WHEN APPLIED TO MORPION

Name	Search Component	Rank	5T	5D	5T, B = 10 ⁵
Dis#1	step(select(step(simulate),0.5))	1	91.43	63.57	97.28
Dis#7	lookahead(step(step(select(simulate, 0))))	2	91.30	63.66	96.28
Dis#4	step(select(step(select(simulate,0.5)),0))	3	91.28	63.61	96.23
Dis#3	step(select(step(select(simulate,1.0)),0))	4	91.27	63.61	96.13
Dis#8	step(select(step(step(simulate)),1))	5	91.19	63.68	96.51
Dis#2	step(step(select(simulate,0)))	6	91.18	63.63	96.62
Dis#9	step(select(step(select(simulate,0)),0.3))	7	91.17	63.67	95.99
Dis#5	select(step(select(step(simulate),1.0)),0)	8	91.16	63.74	95.86
Dis#10	step(select(step(select(simulate,1.0)),0.0))	9	91.12	63.61	95.98
Dis#6	lookahead(step(step(simulate)))	10	91.01	63.84	96.38
nmc(4)		11	90.68	63.36	95.43
nmc(3)		12	90.67	63.51	95.61
la(1)		13	90.63	63.43	95.31
nmc(2)		14	90.54	<u>63.52</u>	95.50
nmc(5)		15	90.44	63.46	95.31
uct(0)		16	90.35	63.02	92.79
uct(0.5)		17	90.25	62.91	92.33
uct(1)		18	90.22	63.12	92.80
uct(0.3)		19	90.01	63.03	92.54
la(2)		20	88.91	62.67	<u>95.76</u>
la(3)		21	85.87	61.52	89.35
is		21	85.87	61.40	89.15
la(4)		23	85.80	61.53	88.87
mcts		24	85.79	61.48	89.15
la(5)		25	85.78	61.52	88.87

and [33] evaluates different versions of a two-player MCS algorithm on generic sparse bandit problems. [34] provides an in-depth review of different MCS algorithms and their successes in different applications.

The main feature of the approach proposed in the present paper is that it builds the space of candidate algorithms by using a rich grammar over the search components. In this sense, [35], [36] are certainly the papers which are the closest to ours, since they also use a grammar to define a search space, for, respectively, two player games and multi-armed bandit problems. However, in both cases, this grammar only models a selection policy and is made of classic functions such as +, −, *, /, log, exp, and $\sqrt{\cdot}$. We have taken one step forward, by directly defining a grammar over the MCS algorithms that covers very different MCS techniques. Note that the search technique of [35] is based on genetic programming.

The decision as to what to use as the performance criterion is not as trivial as it looks, especially for multi-player games, where opponent modelling is crucial for improving over game-theoretically optimal play [37]. For example, the maximization of the victory rate or loss minimization against a wide variety of opponents for a specific game can lead to very different choices of algorithms. Other examples of criteria to discriminate between algorithms are simple regret [29] and the expected performance over a distribution density [38].

VII. CONCLUSION

In this paper we have addressed the problem of automatically identifying new Monte Carlo search (MCS) algorithms performing well on a distribution of training problems. To do so, we introduced a grammar over the MCS algorithms that generates a very rich space of candidate algorithms (and which describes, along the way, using a particularly compact

and elegant description, several well known MCS algorithms). To efficiently search inside this space of candidate algorithms for the one(s) having the best average performance on the training problems, we relied on a multi-armed bandit type of optimisation algorithm.

Our approach was tested on three different domains: Sudoku, Morpion Solitaire, and Symbolic Regression. The results showed that the algorithms discovered this way can significantly outperform generic algorithms such as UCT or NMC on each of these domains. Moreover, we showed that they had good robustness properties, by changing the testing budget and/or by using a testing problem distribution different from the training distribution.

This work can be extended in several ways. For the time being, we used the mean performance over a set of training problems to discriminate between different candidate algorithms. One direction for future work would be to adapt our general approach to use other criteria, e.g., worst case performance measures. In its current form, our grammar only allows using predefined simulation policies. Since the simulation policy typically has a major impact on the performance of a MCS algorithm, it could be interesting to extend our grammar so that it could also “generate” new simulation policies. This could be arranged by adding a set of simulation policy generators in the spirit of our current search component generators. Of course, working with richer grammars will lead to larger candidate algorithm spaces, which in turn, may require developing more efficient search methods than the multi-armed bandit one used in this paper. Finally, another important direction for future research is to extend our approach to more general settings than single-player games with full observability.

ACKNOWLEDGMENT

This paper presents research results of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office. The scientific responsibility rests with its author(s).

REFERENCES

- [1] L. Kocsis and C. Szepesvári, “Bandit based Monte Carlo planning,” *Proceedings of the 17th European Conference on Machine Learning (ECML 2006)*, pp. 282–293, 2006.
- [2] R. Coulom, “Efficient selectivity and backup operators in Monte Carlo tree search,” in *Proceedings of the 5th International Conference on Computers and Games, Turin, Italy*, 2006.
- [3] T. Cazenave, “Nested Monte Carlo search,” in *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, 2009, pp. 456–461.
- [4] J. Méhat and T. Cazenave, “Combining UCT and Nested Monte Carlo Search for Single-Player general game playing,” *IEEE Trans. Comput. Intellig. and AI in Games*, vol. 2, no. 4, pp. 271–277, 2010.
- [5] P. Auer, P. Fischer, and N. Cesa-Bianchi, “Finite-time analysis of the multi-armed bandit problem,” *Machine Learning*, vol. 47, pp. 235–256, 2002.
- [6] G. Tesauro and G. R. Galperin, “On-line policy improvement using Monte Carlo Search,” in *Advances in Neural Information Processing Systems 9, NIPS, 1996*, pp. 1068–1074.
- [7] T. Cazenave, “Reflexive Monte Carlo search,” in *Computer Games Workshop 2007 (CGW 2007)*, Amsterdam, 2007, pp. 165–173.
- [8] G.-B. Chaslot, J.-B. Hoock, J. Perez, A. Rimmel, O. Teytaud, and M. Winands, “Meta Monte Carlo tree search for automatic opening book generation,” in *IJCAI 2009 Workshop on General Intelligence in Game Playing Agents*, 2009, pp. 7–12.
- [9] F. Maes, L. Wehenkel, and D. Ernst, “Learning to play K -armed bandit problems,” in *International Conference on Agents and Artificial Intelligence (ICAART 2012)*, Vilamoura, Algarve, Portugal, February 2012.
- [10] —, “Meta-learning of exploration/exploitation strategies: The multi-armed bandit case,” in *ICAART 2012 - Springer Selection*, arXiv:1207.5208, 2012.
- [11] M. Castronovo, F. Maes, R. Fonteneau, and D. Ernst, “Learning exploration/exploitation strategies for single trajectory reinforcement learning,” in *10th European Workshop on Reinforcement Learning (EWRL 2012)*, Edinburgh, Scotland, June 2012.
- [12] F. Maes, R. Fonteneau, L. Wehenkel, and D. Ernst, “Policy search in a space of simple closed-form formulas: Towards interpretability of reinforcement learning,” in *Discovery Science (DS 2012)*, Lyon, France, October 2012.
- [13] S. Bubeck, R. Munos, and G. Stoltz, “Pure exploration in multi-armed bandits problems,” in *Algorithmic Learning Theory*, 2009, pp. 23–37.
- [14] J. Koza and R. Poli, “Genetic programming,” in *Search Methodologies*, E. K. Burke and G. Kendall, Eds. Berlin: Springer-Verlag, 2005, pp. 127–164.
- [15] P. Auer, N. Cesa-Bianchi, and P. Fischer, “Finite-time analysis of the multiarmed bandit problem,” *Machine learning*, vol. 47, no. 2, pp. 235–256, 2002.
- [16] P.-A. Coquelin and R. Munos, “Bandit algorithms for tree search,” in *Conference on Uncertainty in Artificial Intelligence (UAI 2007)*, Vancouver, Canada, 2007.
- [17] J.-B. Hoock and O. Teytaud, “Bandit-based genetic programming,” in *Proceedings of the 13th European Conference on Genetic Programming*, ser. EuroGP’10. Berlin: Springer-Verlag, 2010, pp. 268–277. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-12148-7_23
- [18] Z. Geem, “Harmony search algorithm for solving sudoku,” in *International Conference on Knowledge-Based Intelligent Information and Engineering Systems (KES 2007)*. Berlin: Springer-Verlag, 2007, pp. 371–378.
- [19] T. Cazenave, “Nested Monte Carlo expression discovery,” in *Proceedings of the European Conference in Artificial Intelligence (ECAI 2010)*, Lisbon, 2010, pp. 1057–1058.
- [20] N. Uy, N. Hoai, M. O’Neill, R. McKay, and E. Galván-López, “Semantically-based crossover in genetic programming: Application to real-valued symbolic regression,” *Genetic Programming and Evolvable Machines*, vol. 12, no. 2, pp. 91–119, 2011.
- [21] C. Boyer, “<http://www.morpionsolitaire.com>,” 2012.
- [22] C. Rosin, “Nested rollout policy adaptation for Monte Carlo tree search,” in *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011)*, Barcelona, Spain, 2011, pp. 649–654.
- [23] E. Demaine, M. Demaine, A. Langerman, and S. Langerman, “Morpion solitaire,” *Theory of Computing Systems*, vol. 39, no. 3, pp. 439–453, 2006.
- [24] P. Perrick, D. St-Pierre, F. Maes, and D. Ernst, “Comparison of different selection strategies in Monte Carlo tree search for the game of Tron,” in *Proceedings of the IEEE Conference on Computational Intelligence and Games (CIG 2012)*, Granada, Spain, 2012.
- [25] I. Chaslot, M. Winands, and H. van den Herik, “Parameter tuning by the cross-entropy method,” in *Proceedings of the European Workshop on Reinforcement Learning (EWRL 2008)*, 2008.
- [26] R. Coulom, “CLOP: Confident local optimization for noisy black-box parameter tuning,” in *Proceedings of the 13th International Advances in Computer Games Conference (ACG 2011)*, 2011.
- [27] F. Maes, L. Wehenkel, and D. Ernst, “Optimized look-ahead tree search policies,” in *Proceedings of the 9th European workshop on reinforcement learning (EWRL’11)*, Athens, Greece, September 2011.
- [28] O. Chapelle and L. Li, “An empirical evaluation of Thompson sampling,” in *Neural Information Processing Systems (NIPS)*, 2011.
- [29] A. Bourki, M. Coulom, P. Rolet, O. Teytaud, P. Vayssi  re et al., “Parameter tuning by simple regret algorithms and multiple simultaneous hypothesis testing,” in *Proceedings of the International Conference on Informatics in Control, Automation and Robotics (ICINCO 2010)*, 2010.
- [30] V. Berthier, H. Daghmen, and O. Teytaud, “Consistency modifications for automatically tuned Monte Carlo tree search,” *Proceedings of the 4th Conference on Learning and Intelligent Optimization*, pp. 111–124, 2010.
- [31] S. Gelly and D. Silver, “Combining online and offline knowledge in UCT,” in *Proceedings of the 24th International Conference on Machine Learning (ICML 2007)*, 2007, pp. 273–280.
- [32] G. Chaslot, M. Winands, H. Herik, J. Uiterwijk, and B. Bouzy, “Progressive strategies for Monte Carlo tree search,” in *Proceedings of the 10th Joint Conference on Information Sciences (JCIS 2007)*, 2007, pp. 655–661.
- [33] D. St-Pierre, Q. Louveaux, and O. Teytaud, “Online sparse bandit for card game,” in *Proceedings of Advances in Computer Games 2011 (ACG 2011)*, 2011.
- [34] C. Browne, E. Powley, D. Whitehouse, S. Lucas, P. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton, “A survey of Monte Carlo tree search methods,” *IEEE Transactions on Computational Intelligence and AI in Games*, no. 99.
- [35] T. Cazenave, “Evolving Monte Carlo tree search algorithms,” Tech. Rep., 2007.
- [36] F. Maes, L. Wehenkel, and D. Ernst, “Automatic discovery of ranking formulas for playing with multi-armed bandits,” in *Proceedings of the 9th European Workshop on Reinforcement Learning (EWRL 2011)*, 2011.
- [37] D. Billings, A. Davidson, J. Schaeffer, and D. Szafron, “The challenge of poker,” *Artificial Intelligence*, vol. 134, no. 1-2, pp. 201–240, 2002.
- [38] V. Nannen and A. Eiben, “Relevance estimation and value calibration of evolutionary algorithm parameters,” in *Proceedings of the 20th International Joint Conference on Artificial Intelligence*. Morgan Kaufmann Publishers, 2007, pp. 975–980.